CONCEPT WHITEPAPER

A Municipal Boiler That Happens to Think

Industrial-first thermal-compute districts for data-center heat reuse, water resilience, and urban expansion

CENTRAL THESIS

A data center should be planned not only as an electrical and computational load, but as a steady thermal utility. Prove the model in an industrial district, where land, substations, pipes, tanks, and continuous heat customers already exist; then extend the thermal network toward hospitals, housing, hotels, public baths, food systems, and smaller urban compute nodes.

Prepared for Nathan Ribeiro

Version 1.0 | 12 July 2026

Concept-stage document. Site engineering, water balance, grid studies, permitting, and commercial diligence remain mandatory.

Abstract

AI and accelerated computing are turning data centers into large, continuous electrical loads. Nearly all of the electricity entering the IT equipment ultimately leaves as heat. The conventional facility treats that heat as a liability to be rejected through air, refrigeration, dry coolers, or evaporation. This whitepaper proposes a different planning model: treat compute as the anchor of a local thermal utility, while preserving independent cooling and uptime.

The proposed deployment sequence begins on the industrial edge of a city, where land, high-voltage access, fiber, water infrastructure, process-heat consumers, and tolerance for heavy equipment are more readily available. A liquid-cooled compute plant feeds an isolated thermal hub containing heat exchangers, heat pumps, storage, metering, and backup rejection. Multiple customers consume the heat in a temperature cascade: sanitation and process water first, then domestic hot water and pools, followed by greenhouses, aquaculture, drying preheat, or wastewater treatment. Once the industrial system has operating history and standardized contracts, its thermal spine can extend inward and support smaller urban compute nodes.

The concept is not a perpetual-energy loop, and low-temperature heat is generally poor feedstock for regenerating electricity. Its value lies in replacing heat that nearby users would otherwise purchase, reducing evaporative water consumption where dry rejection is practical, and converting a sealed data-center project into shared civic infrastructure.

Contents

1. The proposition

THE IDEA IN ONE SENTENCE

Build the data center as a municipal boiler that happens to think.

A data center is commonly evaluated as a consumer of electricity, land, grid capacity, water, and public tolerance. Thermally, however, it is also a remarkably predictable producer: a facility operating continuously converts almost all of its electrical input into heat. The industry has long understood heat recovery, but most projects still optimize the data center as an isolated building rather than as one organ in a district energy system. [2][4]

The planning inversion proposed here is simple: cooling is not merely the act of deleting heat. Cooling is the controlled transfer of heat from compute equipment to a useful destination. The project should therefore be designed around both sides of the transaction: a source with strict uptime requirements and a portfolio of heat users with different temperature, volume, and availability needs.

The initial deployment does not belong in the densest urban core. It belongs at the industrial-urban seam: close enough to heat users and future city branches, but with adequate land, substations, logistics, noise tolerance, security setbacks, and pipe routes. The first project proves the thermal and commercial system. Later projects can move closer to dense demand.

Design principles

- Computation remains the primary mission; heat sales may never compromise thermal safety or uptime.
- Rack coolant, facility water, district water, potable water, and process water remain physically isolated through heat exchangers.
- Heat is cascaded from higher-value temperature requirements to lower-value uses before final rejection.
- No project depends on one greenhouse, one factory, or one property owner. Demand is diversified.
- The system retains full independent heat-rejection capacity for outages, maintenance, and demand mismatch.
- Thermal storage buffers hours and daily peaks; it does not magically solve every seasonal imbalance.
- Expansion is modular: industrial pilot, thermal spine, district customers, then urban compute nodes.

2. Why now

The opportunity is becoming larger and more geographically concentrated. The International Energy Agency projects global data-center electricity consumption to rise from about 485 TWh in 2025 to roughly 950 TWh in 2030, with AI-focused facilities growing considerably faster than the sector as a whole. [1] The local effect is more important than the global percentage: large projects can resemble power-intensive factories and can dominate new demand in particular grid regions.

At the same time, accelerated computing is pushing rack density beyond the comfortable range of conventional air cooling. Direct liquid cooling moves heat into water or another liquid at a higher and more useful temperature, avoiding one air-to-water transfer and making dry rejection or heat reuse more practical. [3][4][5]

Water scrutiny is also rising. In evaporative systems, water is intentionally consumed as heat is rejected, with additional blowdown required to control mineral concentration. [7] Warm-water liquid cooling combined with dry coolers can materially reduce operational cooling-water use in suitable climates, although no universal zero-water claim is credible without a site-specific design and climate analysis. [3][5]

Finally, policy is beginning to treat energy and water performance as reportable infrastructure outcomes. The European Union now requires qualifying data centers to report energy and water indicators, and its framework explicitly recognizes energy reuse. [11] Even where regulation is weaker, community opposition, grid constraints, water permits, and sustainability commitments create the same practical pressure: new capacity needs a better civic bargain.

3. Thermodynamic foundation

The concept begins with a distinction between energy and useful work. A large data center may produce tens or hundreds of megawatts of heat, yet the coolant may be only moderately hotter than its environment. That heat has low exergy: it is valuable for another low-temperature heat requirement, but inefficient to convert back into electricity.

The correct target is therefore substitution, not regeneration. A hospital, hotel, dairy, laundry, digester, pool, greenhouse, or drying process may already purchase electricity, gas, steam, or biomass to raise water from ambient temperature. Datacenter heat can provide part or all of that lift directly, or a heat pump can upgrade it to a higher delivery temperature. Heat pumps are particularly useful with low-temperature sources below approximately 45°C and can connect them to district systems. [12]

Temperature cascading improves utilization. A high-temperature customer accepts the hottest supply and returns cooler water; a second customer uses the return; a third consumes the remaining low-grade heat. This does not reuse the same energy without limit. It extracts useful service from successive temperature drops before the final sink.

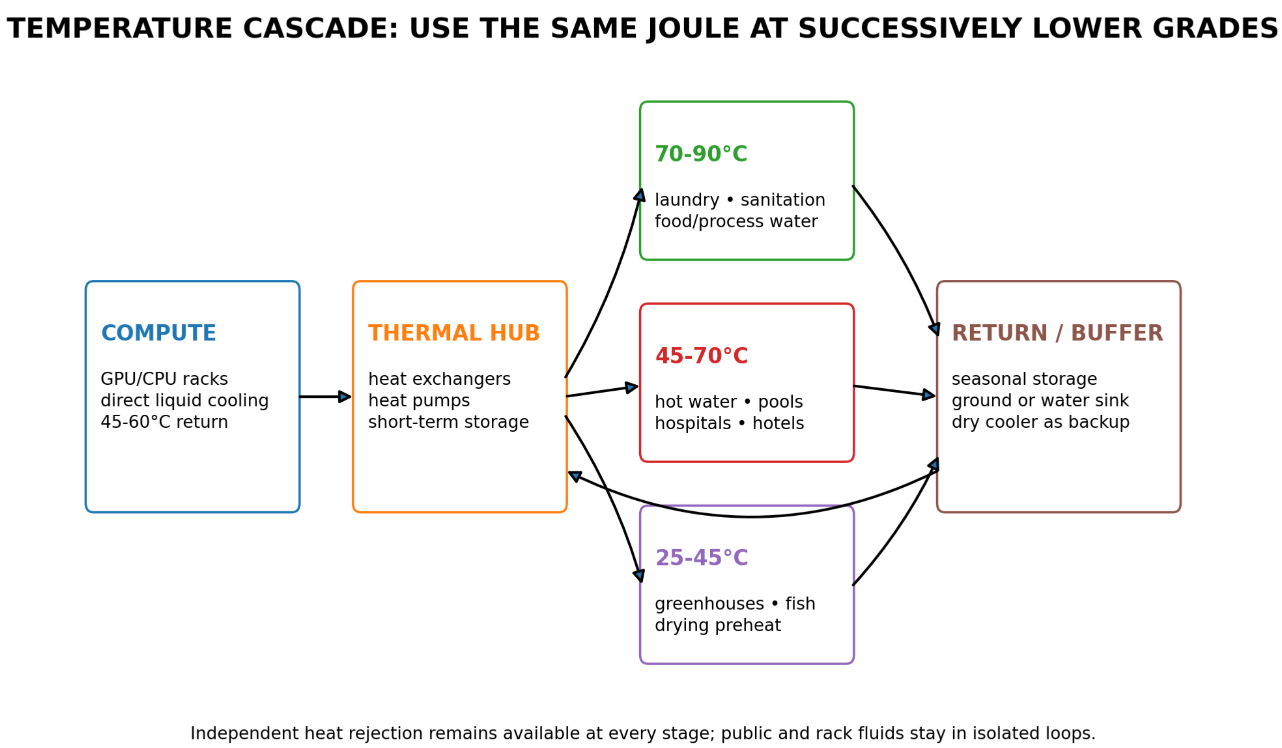

Figure 1. A temperature cascade turns heat rejection into a sequence of useful thermal services.

What the proposition is not

- It is not a method for recovering most of the electricity consumed by the chips.
- It is not an open bath of server coolant; the loops remain isolated.
- It is not permission to remove independent cooling systems.
- It is not a claim that every climate or every datacenter has a viable heat market.
- It is not a single-purpose greenhouse bolted onto a hyperscale campus.

4. Reference system architecture

The reference architecture separates responsibilities into five physical layers. This separation is important because the compute operator, thermal utility, and heat customers have different safety, chemistry, pressure, maintenance, and uptime requirements.

Layer	Function
1. IT thermal capture	Cold plates or immersion systems remove heat from processors and memory. A cooling distribution unit isolates the technology loop from facility water.
2. Facility loop	Pumps move warm water between the data hall and the thermal hub. The design maximizes leaving-water temperature while remaining inside equipment limits.
3. Thermal hub	Plate heat exchangers, heat pumps, valves, meters, controls, and short-term storage transform a variable source into contracted heat products.
4. District thermal bus	Insulated supply and return pipes connect multiple producers, storage assets, and customers at defined temperatures and pressures.
5. Residual rejection	Dry coolers, chillers, cooling towers, or another independent sink reject heat whenever the network cannot accept it.

Four-loop separation

A robust project normally contains at least four distinct hydraulic domains: the rack or technology loop, the facility loop, the district energy loop, and each customer’s process or building loop. Heat crosses the boundaries; fluid does not. Cooling distribution units and heat exchangers limit contamination and leak propagation and simplify maintenance. [13]

Control philosophy

The control system optimizes for safe compute temperatures first, contracted thermal delivery second, and efficiency third. It forecasts compute load, customer demand, weather, storage state, and electricity prices. When demand is insufficient, the facility diverts heat to storage or rejection without asking workloads to wait. Flexible workloads may provide additional optimization, but should be treated as optional value rather than the foundation of cooling safety.

5. Industrial-first deployment strategy

The industrial district is the correct first habitat because the physical and commercial pieces already exist in partial form: high-capacity electricity, industrial water, fiber corridors, process heat demand, large parcels, logistics access, and operators accustomed to utility contracts. Construction can be greenfield, allowing pipe corridors and tanks to be designed before roads, buildings, and property boundaries harden around them.

The pilot should not begin with a speculative heat customer. It should begin with a heat-demand map and two or more anchor users whose combined demand is continuous, contractible, and temperature-compatible. Suitable anchors include industrial laundries, hospitals, hotel or residential hot-water systems, food and beverage plants, dairy and slaughterhouse sanitation, wastewater digesters, timber or agricultural drying preheat, aquaculture, and pool or bathing complexes.

Deployment gates

1. Map hourly and seasonal heat demand within practical pipe distances, including return temperatures.
2. Select IT cooling technology and facility-water temperatures before promising a heat product.
3. Design the temperature cascade and identify which customers receive first-pass and return heat.
4. Price the entire avoided system: cooling plant, water, customer fuel, carbon, pipes, pumps, and backup capacity.
5. Secure long-term rights-of-way and thermal offtake contracts before final datacenter siting.
6. Preserve full cooling redundancy and define failure responsibility across organizational boundaries.
7. Build metering and settlement rules at the start, not after commissioning.
8. Phase capacity so that compute and heat demand can grow without stranding the network.

6. Urban expansion model

After the industrial node operates successfully, the most valuable asset is not the greenhouse or the datacenter itself. It is the trusted thermal utility: standardized interfaces, known return temperatures, operating data, tariffs, rights-of-way, backup rules, and a portfolio of customers. That root system can extend toward the city.

The urban strategy should be incremental. A branch first reaches a hospital, hotel cluster, university, public pool, housing development, food market, or wastewater facility. Thermal density rises along the corridor. Smaller edge data centers can then connect as additional heat producers near demand, while also reducing latency and supporting local compute resilience. NREL has identified proximity to end users and heat-reuse opportunities as potential advantages of distributed edge data centers. [10]

This is fundamentally different from placing a hyperscale campus in a congested downtown parcel. Urban nodes can be smaller, quieter, more architecturally integrated, and connected to an established thermal system whose commercial and safety behavior is already understood.

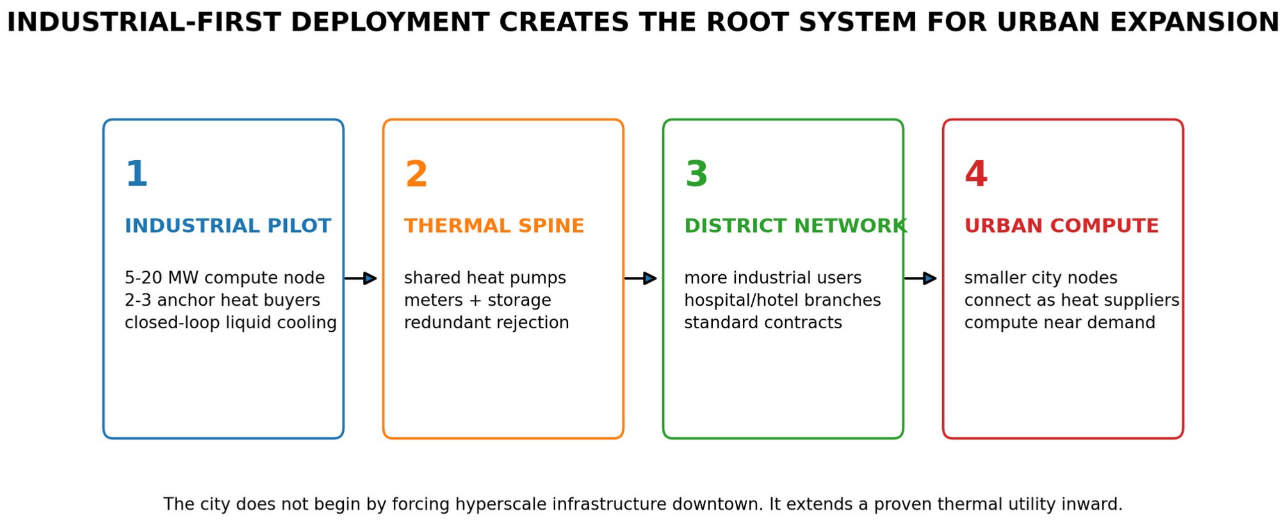

Figure 2. Industrial deployment creates the technical and contractual root system for later urban nodes.

7. Water, cooling, and storage

The public debate often combines several different water flows. The rack and facility loops are closed systems that transport heat repeatedly, aside from maintenance losses. Cooling-tower water is different: it is consumed through evaporation and blowdown as the facility rejects heat. Potable, sanitary, agricultural, and industrial process water are different again and must not be conflated with server coolant.

Heat reuse can reduce the amount of heat that must be deleted outdoors. Warm-water designs may also allow dry coolers to replace evaporative rejection for more hours of the year. The water benefit therefore comes from system architecture, not from claiming that heat recovery itself creates water.

Water domain	Typical status	Purpose	Key rule
Rack / technology loop	Closed	Processor heat transport	Isolated by CDU; chemistry controlled
Facility loop	Closed	Moves heat to thermal hub	Designed for warm return and low pumping energy
District thermal loop	Closed	Sells and redistributes heat	Utility-owned or concessioned; metered
Cooling tower makeup	Consumed	Residual heat rejection	Evaporation plus blowdown; minimize where feasible
Customer process water	Varies	Laundry, sanitation, food, bathing	Receives heat through exchanger; never mixed with coolant

Thermal storage

Hot-water tanks are mature and useful for minute-to-hour balancing, peak shaving, maintenance windows, and morning or evening demand peaks. Larger tanks or pit storage can cover longer periods where land is available. Borehole or aquifer storage may support seasonal strategies, but changes the project into a substantial district-energy development and requires geotechnical and regulatory diligence.

Storage reduces simultaneity requirements; it does not remove the need for a viable annual heat market. A warm-climate project should favor year-round uses such as sanitation, hot water, drying, digestion, aquaculture, and absorption or adsorption cooling rather than depending mainly on winter space heating.

8. Commercial and institutional structure

The thermal network should be treated as a utility rather than as an informal sustainability attachment. The cleanest structure separates the compute operator from the thermal-system owner, while aligning incentives through service contracts.

Party	Role	Commercial logic
Compute operator	Produces heat and requires guaranteed cooling	Pays cooling-service fee or receives heat credit; retains emergency rejection
Thermal utility	Owns hub, pumps, heat pumps, storage, pipes, and metering	Buys or accepts source heat; sells defined thermal products
Anchor customers	Consume contracted heat with defined return temperature	Pay per delivered MWh_th and availability class
Municipality / development authority	Coordinates zoning, rights-of-way, public anchors, and resilience	Grants concessions, aggregates demand, sets disclosure requirements
Grid / energy supplier	Provides electricity and potentially demand-flexibility value	Coordinates connection, tariffs, storage, and flexible load where appropriate

The business case must count avoided systems

Heat should not be evaluated only against the marginal price of fuel. A credible model includes avoided or downsized cooling equipment, reduced evaporative water treatment, deferred customer boilers, avoided customer electrical demand, emissions obligations, maintenance, network losses, backup systems, and financing. Conversely, it must include pipes, rights-of-way, pumping, heat pumps, storage, metering, interconnection, and contractual risk.

Stockholm's Open District Heating model demonstrates one market design: businesses with excess heat can sell it into an established network. [9] The lesson is not that every city can copy Stockholm. It is that heat becomes bankable when there is a buyer of last resort, published technical requirements, metering, and a network larger than any single bilateral relationship.

Contract essentials

- Supply and return temperatures, flow, pressure, and water chemistry.
- Minimum and maximum thermal capacity by hour and season.
- Availability classes and planned-maintenance windows.
- Priority rules during compute outages, customer outages, and grid events.
- Heat-pump electricity ownership and exposure to power prices.
- Metering boundary, network losses, and settlement method.
- Change-of-control, customer closure, and stranded-asset provisions.
- Carbon-accounting method and prohibition on double counting.

9. Reference pilot: a 20 MW node

The following sizing is illustrative, not a design promise. It shows the order of magnitude and the questions a feasibility study must answer. A 20 MW continuous IT load produces approximately 480 MWh of heat per day at the servers. The IEA estimates that roughly 70-80% of data-center electricity consumption can be recoverable as useful heat with heat pumps, depending on system design and demand. [2]

Parameter	Illustrative value	Purpose / caveat
IT load	20 MW continuous	Reference scale for an industrial pilot
Raw server heat	~480 MWh_th/day	Approximately equal to IT electricity over 24 hours
Potential useful recovery	~14-16 MW / 336-384 MWh_th/day	Illustrative 70-80% range; site-specific
Initial customer portfolio	2-3 anchor users plus one flexible sink	Avoid single-customer dependency
Thermal storage	2-6 hours of contracted delivery	Buffers peaks and short interruptions
Pipe radius	Preferably <5 km for first phase	Actual economics depend on thermal density and rights-of-way
Independent rejection	100% of compute heat	Mandatory for uptime and commercial independence
Expansion reserve	Valves, pumps, and route for additional branches	Build the root system without overbuilding all capacity

Example customer stack

Temperature tier	Example user	System value
First pass: 75-90°C after heat pump	Hospital laundry + sterilization preheat	Steady, high-value sanitation heat

Temperature tier	Example user	System value
Second pass: 50-70°C	Hotel/residential hot water + public baths	Daily diversity and civic visibility
Third pass: 30-45°C	Digester, aquaculture, greenhouse root heating	Flexible low-temperature sink
Final buffer / rejection	Tank, ground store, dry cooler, tower	Guarantees compute independence

Feasibility-study outputs

- Hourly source and demand curves for at least one representative year.
- Hydraulic and temperature model including return-water behavior.
- Heat-pump COP map across weather and delivery temperatures.
- Water balance for all normal and emergency cooling modes.
- Electrical one-line impact, grid connection, and backup-generation plan.
- Pipe route, land ownership, geotechnical risks, and permitting schedule.
- Levelized cost of delivered heat with sensitivity to electricity and customer attrition.
- Failure-mode and effects analysis spanning datacenter and utility boundaries.

10. Risks and design responses

Risk	Design response
Heat demand is seasonal or smaller than expected	Diversify year-round users; stage compute capacity; retain dry rejection; use storage for short mismatch only
Customer closes or breaches contract	Multiple offtakers; credit support; utility as aggregator; modular network branches
Datacenter load falls or facility is offline	Backup boilers or electric heat at customer side; storage; no critical customer relies on one compute site
Heat is too cold for the intended process	Increase source temperature where equipment permits; add heat pumps; assign lower-temperature users
Heat-pump electricity erodes economics	Optimize lift and return temperature; use cascading; secure tariffs; compare with customer fuel and avoided cooling
Leaks or water chemistry threaten IT	CDUs and multiple isolated loops; leak detection; pressure hierarchy; maintainability and bypasses
Urban construction becomes politically or physically difficult	Begin on industrial land; secure corridors early; expand with smaller branches and nodes
Sustainability claims overstate benefits	Meter delivered heat and water; use conservative baselines; disclose network losses and backup energy

Risk	Design response
Thermal utility interferes with uptime	Compute-first controls; full residual rejection; contractual demarcation; routine failover testing

NON-NEGOTIABLE RELIABILITY RULE

The heat network must be valuable enough to use, but never necessary for the datacenter to survive. Likewise, critical heat customers must have backup capacity independent of any single datacenter.

11. Measurement and verification

PUE alone is insufficient because it measures facility overhead but does not reward useful heat delivery. The project should report energy, water, thermal delivery, and service reliability together. DOE guidance explicitly recommends heat reuse and the Energy Reuse Effectiveness metric, while European reporting rules include energy and water indicators. [6][11]

Metric	Definition	Why it matters
PUE	Total facility energy divided by IT energy	Cooling and electrical overhead; not a reuse metric
WUE	Operational water use normalized to IT energy	Cooling-water burden; disclose source and boundary
ERF / ERE	Fraction reused / effectiveness after reuse	Whether exported energy is genuinely useful
Delivered heat	MWh_th delivered by temperature tier	Primary product of the thermal utility
Recovery rate	Useful exported heat / available source heat	Technical utilization, reported hourly and annually
Heat-pump performance	COP by source and delivery temperature	Cost of upgrading low-grade heat
Thermal availability	Contracted delivery met over time	Utility reliability
Residual rejection	MWh_th rejected by dry, chiller, or tower mode	Shows remaining heat and water deletion
Avoided energy and emissions	Measured customer baseline minus replacement energy	Must include backup energy and network losses

Public reporting

Because large datacenters can affect grid capacity, water permits, local heat markets, and public infrastructure, municipalities should require site-level disclosure of electricity, operational water, heat exported, heat rejected, backup-generator operation, and verified customer use. Aggregate corporate sustainability reports are not a substitute for local infrastructure accounting.

12. Policy and planning agenda

The concept requires coordination that ordinary parcel-by-parcel permitting does not provide. The municipality or development authority does not need to operate the datacenter, but it must make the thermal market legible and investable.

- Create a geospatial heat-demand atlas including hospitals, laundries, hotels, wastewater plants, food processors, pools, and planned housing.
- Require large datacenter proposals to submit a heat-reuse and water-rejection feasibility study before siting approval.
- Reserve utility corridors for insulated thermal pipes alongside power, water, sewer, and fiber.
- Use public facilities as anchor customers where lifecycle economics are favorable.
- Permit a thermal-utility concession or open-access network with published technical connection rules.
- Reward high-temperature liquid cooling, verified heat export, and low operational water use without prescribing one vendor technology.
- Require full backup and emergency rejection so public heat contracts cannot compromise digital resilience.
- Coordinate grid connection, compute flexibility, heat-pump load, and thermal storage rather than evaluating them as separate projects.
- Phase urban densification around proven industrial infrastructure instead of forcing the first experiment into the city center.

Warm-climate adaptation

Cold-climate district heating is the most visible precedent, but the industrial-first model is not limited to cold cities. Warm regions should prioritize year-round sanitary hot water, food and beverage processes, laundries, wastewater digestion, low-temperature drying, aquaculture, pools and baths, and heat-driven cooling where technically and economically justified. The result may look less like a Nordic heating grid and more like an industrial hot-water utility with urban branches.

13. Conclusion

THE STRATEGIC SEQUENCE

Industrial compute plant → thermal hub → diversified industrial heat market → shared thermal spine → civic customers → smaller urban compute nodes.

The main innovation is not a new heat exchanger. It is a different order of operations. Today, cities usually approve a datacenter, then ask whether some of its waste heat can be donated to a nearby use. The proposed model begins by mapping heat demand, water constraints, electrical capacity, fiber, and urban growth together. The compute plant is then placed inside that metabolism.

Industrial land provides the proving ground. It supports the first pipes, tanks, substations, cooling equipment, industrial customers, and operating contracts. Once the system works, the thermal network becomes infrastructure that can move inward. Hospitals, hotels, housing, public baths, laundries, food systems, and wastewater facilities become customers. Smaller urban datacenters become additional producers.

The result is neither a free-energy machine nor a universal answer to data-center impacts. It is a way to stop paying to manufacture useful heat and then paying again to delete it. Done carefully, the data center becomes more than a sealed electrical consumer: it becomes a municipal boiler that happens to think.

References

[1] International Energy Agency (2026). Key Questions on Energy and AI - Executive Summary.

[2] International Energy Agency (2025). Opportunities for district heating in the changing energy landscape.

[3] ASHRAE / PNNL / NEMA (2026). AI Data Center Energy Performance Framework: Energy and Thermal Efficiency.

[4] Uptime Institute (2023). Heat Reuse: A Management Primer.

[5] Uptime Institute (2024). Performance expectations of liquid cooling need a reality check.

[6] U.S. Department of Energy (2024). Best Practices Guide for Energy-Efficient Data Center Design.

[7] U.S. Department of Energy FEMP (2019). Cooling Water Efficiency Opportunities for Federal Data Centers.

[8] National Renewable Energy Laboratory (2014). Energy Performance Testing of Asetek RackCDU System.

[9] Stockholm Exergi. Heat Recovery / Open District Heating.

[10] National Renewable Energy Laboratory (2025). Considerations for Distributed Edge Data Centers and Use Cases.

[11] European Commission. Energy performance of data centres and Regulation (EU) 2024/1364.

[12] International Energy Agency (2020). Heat pumps in district heating and cooling systems.

[13] Uptime Institute (2022). Glossary of Digital Infrastructure Sustainability - Cooling Distribution Unit.

Citation note: numerical examples are illustrative unless explicitly attributed. This paper should be followed by a site-specific heat-demand map, thermodynamic model, water balance, grid study, and commercial feasibility analysis.